UNITED STATES PATENT AND TRADEMARK OFFICE

A

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/927,712	08/10/2001	Takahiro Fujita	16869S-031400US	5769

20350 7590 08/08/2005

TOWNSEND AND TOWNSEND AND CREW, LLP
TWO EMBARCADERO CENTER
EIGHTH FLOOR
SAN FRANCISCO, CA 94111-3834

EXAMINER	
MOFIZ, APU M	

ART UNIT	PAPER NUMBER
2165	

DATE MAILED: 08/08/2005

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No.	Applicant(s)	
	09/927,712	FUJITA ET AL.	
	Examiner	Art Unit	
	Apu M. Mofiz	2165	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *10 June 2005*.

2a)☒ This action is **FINAL**. 2b)☐ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *17-36* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *17-36* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☒ The drawing(s) filed on *10 August 2001* is/are: a)☒ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☐ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☐ Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08) Paper No(s)/Mail Date _____.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____.
5) ☐ Notice of Informal Patent Application (PTO-152)
6) ☐ Other: _____.

DETAILED ACTION

Examiner's Response to Applicant's Remarks

1. Applicant's arguments submitted on 06/10/2005 with respect to claims 17-32 and newly added claims 33-36 have been reconsidered but are not deemed persuasive for the reasons set forth below.

Examiner's Response to Applicant's Remarks are listed below:

2. Applicant argues (under REMARKS section) that, Peters or Baru do not teach a computer system comprising: a first server operative to receive I/O requests from a requesting computer; a second server connected to said first server via a network and operative to receive I/O requests from said requesting computer, a managing computer in data communication with said first and second servers via said network; a first storage unit; and a second storage unit, said first storage unit being accessible by said first server but not by said second Server; said managing computer operative to obtain data access load conditions from said first and second servers, based on said data access load conditions, including a condition in that a load of said first server exceeds a predetermined amount said managing computer operative to: select said second storage unit; copy a first data partition stored in said first storage unit to said second storage unit and subsequently delete said first data partition from said first storage unit; cause said second server to access said second storage unit; and transmit information

to said requesting computer that said first data is to be accessed via said second server.

Examiner respectfully disagrees. Peters teaches **a first server** (i.e., *"Fig.1 illustrates an example computer system 40 in which the present invention may be used. The computer system includes a plurality of storage units 42. A storage unit is a device with a nonvolatile computer-readable medium, such as a disk, on which data may be stored." ... "For example,* ***the storage unit 42 may be a server computer which stores data in a data file in the file system of the server****. There may be an arbitrary number of storage units in the computer system 40. Applications 44 are systems that requests access to the storage units 42 via requests to the storage units over a* ***computer network 46****. The storage units 42 may deliver data to or receive data from the applications 44 over the network 46." ... "Applications 44 also may be called "clients." One or more* ***catalog managers 49*** *may also be used. A* ***catalog manager*** *is a database, accessible by the applications 44, that maintains information about the data available on the storage units 42." ... "data to be stored on the* ***storage units 42*** *is divided into segments. Each segment is copied. Each copy is stored on a different one of the storage units 42." ... "Referring again to FIG. 1, when an application 44 requests access to a selected segment of data on one of the storage units 42, the storage unit places the request on a queue 48 that is maintained for the storage unit. Applications may make such requests independently of each other or any* ***centralized control****, which makes the system more readily scalable. The selection of a storage unit to which a request is sent may be controlled such that random fluctuations in the load applied by multiple applications 44 on multiple storage units 42 are balanced statistically and more equally over all of the storage units 42."* The preceding text excerpts clearly indicate that multiple storage units, which may be **servers** or **storage disks connected to servers**, are linked to a plurality of applications/**clients** through **a centralized control or a catalog manager** (i.e., 49 in Fig.1). The clients and the servers **(i.e., the front-end server/catalog manager or unit 49 in Fig.1 and the back-end servers or unit 40 in Fig.1 are connected through a computer network or unit 46 in Fig. 1. I/O requests from the clients/applications or unit 44 in Fig.1 to the storage units/servers/backend servers or unit 42 are coordinated through management/centralized control/catalog manager or unit 49 in Fig. 1)** are connected through a

computer network. The arbitrary numbers of servers are the **first server, second server, ... , nth server (i.e., unit 40 in Figure 1)**. The first server is associated with the first disk or the first server only can access data from the first disk. The second server is associated with the second disk etc. The clients request i.e. **(I/O requests)** for data from storage units go through the catalog manager. The catalog manager decides to which storage unit or the server the request should go to, is based on the load condition of the storage units. Finally, Applicant's argument that Peters does not teach a first server and a second server, wherein the servers are connected via a network, is not proper.) (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) **operative to receive I/O requests** (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) **from a requesting computer** (i.e., applications or clients or units 44 in Fig.1) (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5); **a second server** (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) **connected to said first server via a network and operative to receive I/O requests** (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) **from said requesting computer** (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5); **a managing computer** (i.e., the centralized control or the catalog manager or unit 49 in Fig. 1) (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) **in data communication with said first and second servers** (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) **via said network** (i.e., unit 46 in Fig. 1); **a first storage unit** (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5); **and a second storage unit** (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5), **said first storage unit being accessible by said first server** (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) **but not by said second server** (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5); **said managing computer** (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) **operative to obtain data access load conditions** (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) **from said first and second servers** (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5), **is based on said data access load conditions, including a condition in that a load of**

said first server exceeds a predetermined amount (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5), **said managing computer** (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) **operative to: select said second storage unit** (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5); **cause said second server to access said second storage unit** (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5); **and transmit information to said requesting computer** (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) **that said first data is to be accessed via said second server** (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5).

Peters does not explicitly teach copying a first data partition stored in said first storage unit to said second storage unit and subsequently delete said first data partition from said first storage unit.

Baru teaches **copying** (i.e., *"The invention provides a method and apparatus for distributing data of a table partitioned across a group of nodes of a parallel database system to achieve substantially uniform resource loading of the nodes while reducing the amount of data movement."* ... *"each of the nodes is associated with storage and processing resources."* ... ***"comparing the resource loading to obtain a substantially uniform distribution with reduced required movement of data and then moving the identified sub-partitions to the nodes have lower resource loading to balance the loading of the node containing partitions of the table."*** The preceding text excerpts clearly indicate that the data is partitioned/divided into multiple nodes/storage units, is based **on the load condition of the nodes/storage units** data is moved between nodes/storage units. **The meaning of "move" by the Microsoft technical dictionary is copying data to the destination and deleting from the originating location**. Therefore, data from one node, which has excessive resource load, would be moved (i.e., copied from one node to the other node and then deleted from the originating node) to the other node with less load and thus the load of the originating node would be reduced. Finally, applicant's argument that Baru does not teach "copying a first data partition stored in said first storage unit to said second storage unit and subsequently delete said first data partition from said first storage unit" is not proper.) (Fig. 2; col 1, lines 60-67; col 2, lines 1-2, lines 25-35) **a first**

data partition (Fig. 2; col 1, lines 60-67; col 2, lines 1-2, lines 25-35) **stored in said first storage unit** (Fig. 2; col 1, lines 60-67; col 2, lines 1-2, lines 25-35) **to said second storage unit** (Fig. 2; col 1, lines 60-67; col 2, lines 1-2, lines 25-35) **and subsequently delete said first data partition from said first storage unit** (Fig. 2; col 1, lines 60-67; col 2, lines 1-2, lines 25-35).

It would have been obvious to a person of ordinary skill in the art at the time of Applicant's invention to modify the teachings of Peters with the teachings of Baru to include copying a first data partition stored in said first storage unit to said second storage unit and subsequently delete said first data partition from said first storage unit with the motivation to achieve a substantially uniform resource loading of the nodes while reducing the amount of data movement (Baru, col 1, lines 60-65).

Any other arguments by the applicant are more limiting than the claimed language.

3. Applicant is inaccurate for the reasons explicitly stated in the First Office Action. Examiner asserts that the Peters, Baru and Microsoft teach Applicant's invention.

4. These reasons have been explicitly stated in the First Office Action.

Claim Rejections - 35 USC § 103

5. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the

invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

6. Claims 17-24, 26-31, 33 and 35-36 are rejected under 35 U.S.C. 103(a) as being unpatentable over Peters et al. (U.S. Patent No. 6,415,373 and Peters hereinafter) in view of Baru et al. (U.S. Patent No. 5,970,495 and Baru hereinafter).

As to claims 17, 26 and 33, Peters teaches a first server (i.e., *"Fig.1 illustrates an example computer system 40 in which the present invention may be used. The computer system includes a plurality of storage units 42. A storage unit is a device with a nonvolatile computer-readable medium, such as a disk, on which data may be stored."* ... *"For example,* ***the storage unit 42 may be a server computer which stores data in a data file in the file system of the server****. There may be an arbitrary number of storage units in the computer system 40. Applications 44 are systems that requests access to the storage units 42 via requests to the storage units over a* ***computer network 46****. The storage units 42 may deliver data to or receive data from the applications 44 over the network 46."* ... *"Applications 44 also may be called "clients." One or more* ***catalog managers 49*** *may also be used. A* ***catalog manager*** *is a database, accessible by the applications 44, that maintains information about the data available on the storage units 42."* ... *"data to be stored on the* ***storage units 42*** *is divided into segments. Each segment is copied. Each copy is stored on a different one of the storage units 42."* ... *"Referring again to FIG. 1, when an application 44 requests access to a selected segment of data on one of the storage units 42, the storage unit places the request on a queue 48 that is maintained for the storage unit. Applications may make such requests independently of each other or any* ***centralized control****, which makes the system more readily scalable. The selection of a storage unit to which a request is sent may be controlled such that random fluctuations in the load applied by multiple applications 44 on multiple storage units 42 are balanced statistically and more equally over all of the storage units 42."* The preceding text excerpts clearly indicate that multiple storage units, which may be **servers** or **storage disks connected to servers**, are linked to a plurality of applications/**clients** through **a centralized control or a catalog manager** (i.e., 49 in Fig.1). The clients and the servers **(i.e., the front-end server/catalog manager or**

unit 49 in Fig.1 and the back-end servers or unit 40 in Fig.1 are connected through a computer network or unit 46 in Fig. 1. I/O requests from the clients/applications or unit 44 in Fig.1 to the storage units/servers/backend servers or unit 42 are coordinated through management/centralized control/catalog manager or unit 49 in Fig. 1) are connected through a computer network. The arbitrary numbers of servers are the **first server, second server, ... , nth server (i.e., unit 40 in Figure 1)**. The first server is associated with the first disk or the first server only can access data from the first disk. The second server is associated with the second disk etc. The clients request i.e. **(I/O requests)** for data from storage units go through the catalog manager. The catalog manager decides to which storage unit or the server the request should go to, is based on the load condition of the storage units. Finally, Applicant's argument that Peters does not teach a first server and a second server, wherein the servers are connected via a network, is not proper.) (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) operative to receive I/O requests (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) from a requesting computer (i.e., applications or clients) (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5); a second server (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) connected to said first server via a network and operative to receive I/O requests (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) from said requesting computer (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5); a managing computer (i.e., the centralized control or the catalog manager) (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) in data communication with said first and second servers (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) via said network; a first storage unit (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5); and a second storage unit (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5), said first storage unit being accessible by said first server (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) but not by said second (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) said managing computer (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) operative to obtain data access load conditions (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) from said first and second servers (Fig. 1; col 6, lines 12-67; col 7, lines 59-

67; col 8, lines 1-5),, is based on said data access load conditions, including a condition in that a load of said first server exceeds a predetermined amount (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5), said managing computer (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) operative to: select said second storage unit (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5); cause said second server to access said second storage unit (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5); and transmit information to said requesting computer (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) that said first data is to be accessed via said second server (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5).

Peters does not explicitly teach copying a first data partition stored in said first storage unit to said second storage unit and subsequently delete said first data partition from said first storage unit.

Baru teaches copying (i.e., *"The invention provides a method and apparatus for distributing data of a table partitioned across a group of nodes of a parallel database system to achieve substantially uniform resource loading of the nodes while reducing the amount of data movement."* ... *"each of the nodes is associated with storage and processing resources."* ... ***"comparing the resource loading to obtain a substantially uniform distribution with reduced required movement of data and then moving the identified sub-partitions to the nodes have lower resource loading to balance the loading of the node containing partitions of the table."*** The preceding text excerpts clearly indicate that the data is partitioned/divided into multiple nodes/storage units, is based **on the load condition of the nodes/storage units** data is moved between nodes/storage units. **The meaning of "move" by the Microsoft technical dictionary is copying data to the destination and deleting from the originating location.** Therefore, data from one node, which has excessive resource load, would be moved (i.e., copied from one node to the other node and then deleted from the originating node) to the other node with less load and thus the load of the originating node would be reduced. Finally, applicant's argument that Baru does not teach "copying a first data partition stored in said first storage unit to said second storage unit and subsequently delete said first data partition

from said first storage unit" is not proper.) (Fig. 2; col 1, lines 60-67; col 2, lines 1-2, lines 25-35) a first data partition (Fig. 2; col 1, lines 60-67; col 2, lines 1-2, lines 25-35) stored in said first storage unit (Fig. 2; col 1, lines 60-67; col 2, lines 1-2, lines 25-35) to said second storage unit (Fig. 2; col 1, lines 60-67; col 2, lines 1-2, lines 25-35) and subsequently delete said first data partition from said first storage unit (Fig. 2; col 1, lines 60-67; col 2, lines 1-2, lines 25-35).

It would have been obvious to a person of ordinary skill in the art at the time of Applicant's invention to modify the teachings of Peters with the teachings of Baru to include copying a first data partition stored in said first storage unit to said second storage unit and subsequently delete said first data partition from said first storage unit with the motivation to achieve a substantially uniform resource loading of the nodes while reducing the amount of data movement (Baru, col 1, lines 60-65).

As to claims 18 and 27, Peters teaches that said first data partition is copied (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) from said first storage unit to said second storage unit (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5)

Peters does not explicitly teach that subsequent to said first data partition being copied to said second storage unit, said first data partition is deleted from said first storage unit, thereby effecting a move of said first data partition to said second storage unit.

Baru teaches that subsequent to said first data partition being copied to said second storage unit, said first data partition is deleted from said first storage unit, thereby effecting a move of said first data partition to said second storage unit (i.e., *"The*

invention provides a method and apparatus for distributing data of a table partitioned across a group of nodes of a parallel database system to achieve substantially uniform resource loading of the nodes while reducing the amount of data movement." ... "each of the nodes is associated with storage and processing resources." ... "comparing the resource loading to obtain a substantially uniform distribution with reduced required movement of data and then ***moving*** *the identified sub-partitions to the nodes have lower resource loading to balance the loading of the node containing partitions of the table."* The preceding text excerpts clearly indicate that the data is partitioned/divided into multiple nodes/storage units, is based on the load condition of the nodes/storage unit's data is moved between nodes/storage units. The meaning of "move" by the Microsoft technical dictionary is copying data to the destination and deleting from the originating location.) (Fig. 2; col 1, lines 60-67; col 2, lines 1-2, lines 25-35).

It would have been obvious to a person of ordinary skill in the art at the time of Applicant's invention to modify the teachings of Peters with the teachings of Baru to include that subsequent to said first data partition being copied to said second storage unit, said first data partition is deleted from said first storage unit, thereby effecting a move of said first data partition to said second storage unit with the motivation to achieve a substantially uniform resource loading of the nodes while reducing the amount of data movement (Baru, col 1, lines 60-65).

As to claims 19 and 29, Peters teaches that the said requesting computer is a front-end server that receives requests from client machines, said first server and second server each being a back-end server which receives requests from said front-end server (i.e., *"Fig.1 illustrates an example computer system 40 in which the present invention may be used. The computer system includes a plurality of storage units 42. A storage unit is a device with a nonvolatile computer-readable medium, such as a disk, on which data may be stored." ... "For example, the storage unit 42 may be a server computer which stores data in a data file in the file system of the server. There may be an arbitrary number of*

storage units in the computer system 40. Applications 44 are systems that requests access to the storage units 42 via requests to the storage units over a computer network 46. The storage units 42 may deliver data to or receive data from the applications 44 over the network 46." ... "Applications 44 also may be called "clients." One or more catalog managers 49 may also be used. A catalog manager is a database, accessible by the applications 44, that maintains information about the data available on the storage units 42." ... "data to be stored on the storage units 42 is divided into segments. Each segment is copied. Each copy is stored on a different one of the storage units 42." ... "Referring again to FIG. 1, when an application 44 requests access to a selected segment of data on one of the storage units 42, the storage unit places the request on a queue 48 that is maintained for the storage unit. Applications may make such requests independently of each other or any ***centralized control****, which makes the system more readily scalable. The selection of a storage unit to which a request is sent may be controlled such that random fluctuations in the load applied by multiple applications 44 on multiple storage units 42 are balanced statistically and more equally over all of the storage units 42."* The preceding text excerpts clearly indicate that multiple storage units, which may be **servers** or **storage disks connected to servers**, are linked to a plurality of applications/**clients** through a **centralized control or a catalog manager or the front end server**, which front ends the clients. Therefore the arbitrary numbers of servers are the **first server, second server, ... , nth server**. The first server is associated with the first disk or the first server only can access data from the first disk. The second server is associated with the second disk etc. The clients request i.e. **(I/O requests)** for data from storage units go through the catalog manager. The **catalog manager (i.e., the front end server)** decides to which storage unit or the **server (i.e. the back end servers)** the request should go to, be based on the load condition of the storage units. The data is divided into segments/partitions. Each segment of data is copied to other storage units.) (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5).

As to claims 20 and 23, Peters teaches a storage system, said storage system comprising said first storage unit and said second storage unit, wherein a communication port in said second server can be configured for data communication

with a communication port in said storage system for data access to said second storage unit (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5).

As to claim 21, Peters teaches that a first storage system, comprising said first storage unit and a second storage system comprising said second storage unit (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5).

As to claims 22 and 30, Peters teaches that additional storage units being accessed by additional servers, said requesting computer being informed of said additional servers (i.e., *"Fig.1 illustrates an example computer system 40 in which the present invention may be used. The computer system includes a plurality of storage units 42. A storage unit is a device with a nonvolatile computer-readable medium, such as a disk, on which data may be stored."* ... *"For example, the storage unit 42 may be a server computer which stores data in a data file in the file system of the server. There may be an arbitrary number of storage units in the computer system 40. Applications 44 are systems that requests access to the storage units 42 via requests to the storage units over a computer network 46. The storage units 42 may deliver data to or receive data from the applications 44 over the network 46."* ... *"Applications 44 also may be called "clients." One or more catalog managers 49 may also be used. A catalog manager is a database, accessible by the applications 44, that maintains information about the data available on the storage units 42."* ... *"data to be stored on the storage units 42 is divided into segments. Each segment is copied. Each copy is stored on a different one of the storage units 42."* ... *"Referring again to FIG. 1, when an application 44 requests access to a selected segment of data on one of the storage units 42, the storage unit places the request on a queue 48 that is maintained for the storage unit. Applications may make such requests independently of each other or any* ***centralized control****, which makes the system more readily scalable. The selection of a storage unit to which a request is sent may be controlled such that random fluctuations in the load applied by multiple applications 44 on multiple storage units 42 are balanced statistically and more equally over all of the storage units 42."* The preceding text excerpts clearly indicate that

multiple storage units, which may be **servers** or **storage disks connected to servers**, are linked to a plurality of applications/**clients** through a **centralized control or a catalog manager or the front end server**, which front ends the clients. Therefore the arbitrary numbers of servers are the **first server, second server, ... , nth server**. The first server is associated with the first disk or the first server only can access data from the first disk. The second server is associated with the second disk etc. The clients request i.e. **(I/O requests)** for data from storage units go through the catalog manager. The **catalog manager (i.e., the front end server)** decides to which storage unit or the **server (i.e. the back end servers)** the request should go to, be based on the load condition of the storage units. The data is divided into segments/partitions. Each segment of data is copied to other storage units.) (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5).

Peters does not explicitly teach that moving said first data partition from said first storage unit to said second storage unit includes copying it from said first storage unit to said second storage unit, and subsequent to being copied to said second storage unit said first data partition is deleted from said first storage unit, wherein additional data partitions in said first storage unit can be similarly moved to additional storage units.

Baru teaches that moving said first data partition from said first storage unit to said second storage unit includes copying it from said first storage unit to said second storage unit, and subsequent to being copied to said second storage unit said first data partition is deleted from said first storage unit, wherein additional data partitions in said first storage unit can be similarly moved to additional storage units (i.e., *"The invention provides a method and apparatus for distributing data of a table partitioned across a group of nodes of a parallel database system to achieve substantially uniform resource loading of the nodes while reducing the amount of data movement."* ... *"each of the nodes is associated with storage and processing resources."* ... *"comparing the resource loading to obtain a substantially uniform distribution with reduced required movement of data and then* ***moving*** *the identified sub-partitions to the nodes have lower resource loading to balance the loading of the node containing partitions of the table."* The preceding text excerpts clearly indicate that the data is partitioned/divided

into multiple nodes/storage units, is based on the load condition of the nodes/storage units data is moved between nodes/storage units and that includes 1.. nth storage units. The meaning of "move" by the Microsoft technical dictionary is copying data to the destination and deleting from the originating location.) (Fig. 2; col 1, lines 60-67; col 2, lines 1-2, lines 25-35).

It would have been obvious to a person of ordinary skill in the art at the time of Applicant's invention to modify the teachings of Peters with the teachings of Baru to include that moving said first data partition from said first storage unit to said second storage unit includes copying it from said first storage unit to said second storage unit, and subsequent to being copied to said second storage unit said first data partition is deleted from said first storage unit, wherein additional data partitions in said first storage unit can be similarly moved to additional storage units with the motivation to achieve a substantially uniform resource loading of the nodes while reducing the amount of data movement (Baru, col 1, lines 60-65).

As to claim 24, Peters teaches a switch (i.e., *"In the present invention, topology will vary depending on the physical installation. A non-blocking, switch-based network in which each node, i.e., client or storage unit, is connected directly to the same switch may be used."*) (col 14, lines 28-41) operative for data communication among devices connected to said switch (col 14, lines 28-41), said first server and said second server (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) being connected to said switch (col 14, lines 28-41), said storage system (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) being connected to said switch (col 14, lines 28-41) so that said first and second servers (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) can access data stored in said first and second storage units (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines

1-5), said switch (col 14, lines 28-41) further being operative to direct data requests from one of said first and second servers (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) to one of said first and second storage units (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5).

As to claim 28, Peters teaches that the managing computer is further operative to perform a second configuration operation wherein a second server accesses said second storage unit, if said first configuration operation cannot be performed (i.e., *"Fig.1 illustrates an example computer system 40 in which the present invention may be used. The computer system includes a plurality of storage units 42. A storage unit is a device with a nonvolatile computer-readable medium, such as a disk, on which data may be stored."* ... *"For example, the storage unit 42 may be a server computer which stores data in a data file in the file system of the server. There may be an arbitrary number of storage units in the computer system 40. Applications 44 are systems that requests access to the storage units 42 via requests to the storage units over a computer network 46. The storage units 42 may deliver data to or receive data from the applications 44 over the network 46."* ... *"Applications 44 also may be called "clients." One or more catalog managers 49 may also be used. A catalog manager is a database, accessible by the applications 44, that maintains information about the data available on the storage units 42."* ... *"data to be stored on the storage units 42 is divided into segments. Each segment is copied. Each copy is stored on a different one of the storage units 42."* ... *"Referring again to FIG. 1, when an application 44 requests access to a selected segment of data on one of the storage units 42, the storage unit places the request on a queue 48 that is maintained for the storage unit. Applications may make such requests independently of each other or any* ***centralized control****, which makes the system more readily scalable. The selection of a storage unit to which a request is sent may be controlled such that random fluctuations in the load applied by multiple applications 44 on multiple storage units 42 are balanced statistically and more equally over all of the storage units 42."* The preceding text excerpts clearly indicate that multiple storage units, which may be **servers** or **storage disks connected to servers**, are linked to a plurality of applications/**clients** through a **centralized control or a catalog manager or the front end server**, which front ends

the clients. Therefore the arbitrary numbers of servers are the **first server, second server, ... , nth server**. The first server is associated with the first disk or the first server only can access data from the first disk. The second server is associated with the second disk etc. The clients request i.e. **(I/O requests)** for data from storage units go through the catalog manager. The **catalog manager (i.e., the front end server)** decides to which storage unit or the **server (i.e. the back end servers)** the request should go to, be based on the load condition of the storage units. The data is divided into segments/partitions. Each segment of data is copied to other storage units.) (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5).

As to claim 31, Peters teaches a switching device (col 14, lines 28-41), said first server (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) in data communication with said switching device (col 14, lines 28-41), said managing computer (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) further being operative to obtain loading information (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) from said first server (Fig. 1; col 6, lines 12-67; col 7, lines 59-67; col 8, lines 1-5) or said switching device (col 14, lines 28-41).

As to claims 35 and 36, the limitations of these claims are either addressed or rejected in the rejected claim 17 above.

7. Claim 25,32 and 34 are rejected under 35 U.S.C. 103(a) as being unpatentable over Peters et al. (U.S. Patent No. 6,415,373 and Peters hereinafter) in view of Baru et al. (U.S. Patent No. 5,970,495 and Baru hereinafter) as applied to claims above, and further in view of Microsoft Corporation ("Windows 95/98 PCs- An Introduction", pages 1-4 and Microsoft hereinafter).

As to claims 25,32 and 34, Peters in view of Baru teaches a management computer, a first server, a second server, a first storage unit, a second storage unit, wherein the management computer selects the storage unit in accordance with receiving an indication for moving (i.e., *"Fig.1 illustrates an example computer system 40 in which the present invention may be used. The computer system includes a plurality of storage units 42. A storage unit is a device with a nonvolatile computer-readable medium, such as a disk, on which data may be stored."* ... *"For example, the storage unit 42 may be a server computer which stores data in a data file in the file system of the server. There may be an arbitrary number of storage units in the computer system 40. Applications 44 are systems that requests access to the storage units 42 via requests to the storage units over a computer network 46. The storage units 42 may deliver data to or receive data from the applications 44 over the network 46."* ... *"Applications 44 also may be called "clients." One or more catalog managers 49 may also be used. A catalog manager is a database, accessible by the applications 44, that maintains information about the data available on the storage units 42."* ... *"data to be stored on the storage units 42 is divided into segments. Each segment is copied. Each copy is stored on a different one of the storage units 42."* ... *"Referring again to FIG. 1, when an application 44 requests access to a selected segment of data on one of the storage units 42, the storage unit places the request on a queue 48 that is maintained for the storage unit. Applications may make such requests independently of each other or any* ***centralized control****, which makes the system more readily scalable. The selection of a storage unit to which a request is sent may be controlled such that random fluctuations in the load applied by multiple applications 44 on multiple storage units 42 are balanced statistically and more equally over all of the storage units 42."* The preceding text excerpts clearly indicate that multiple storage units, which may be **servers** or **storage disks connected to servers**, are linked to a plurality of applications/**clients** through a **centralized control or a catalog manager or the front end server**, which front ends the clients. Therefore the arbitrary number of servers is the **first server, second server, ... , nth server**. The first server is associated with the first disk or the first server only can access data from the first disk. The second server is associated with the second disk etc. The clients request i.e. **(I/O requests)** for data from storage units go through the catalog manager. The **catalog manager (i.e., the front end**

server) decides to which storage unit or the **server (i.e. the back end servers)** the request should go to, be based on the load condition of the storage units. The data is divided into segments/partitions. Each segment of data is copied to other storage units.) (Fig. 1; col 6, lines 12-67; col 7, lines 59- 67; col 8, lines 1-5).

Peters in view of Baru does not explicitly teach that the management computer includes a display unit operable to present a first display area and a second display area, said first display area to display one or more first symbols that represent said first server, said second server, or said first storage unit, and having second symbols that represent communication paths, said second display area having third symbols that represent said second storage unit, wherein said management computer selects the second storage unit in accordance with receiving an indication for moving one of said third symbols from said second display area into said first display area.

Microsoft teaches that the management computer (i.e., "A Network Server controls our computer network.") (page 1) includes a display unit operable to present a first display area and a second display area (i.e., *"Mapped drives appear, as an extra drive on your computer, yet are located elsewhere on the network Eg S Drive = Students (Curriculum Server 1) T Drive = Teachers (curriculum Server 1)" ... "Drop and Drag This feature allows you to move files and folders from one folder to another in Windows 95/98"* The preceding text excerpts clearly indicate that a computer in the network e.g., a network server with windows 95/98 shows symbols of other storage units in other servers e.g., curriculum server. A user can move data e.g., a symbol of a file from one server to the other server by dragging the symbol from one area to the other.) (page 1; page 3; page 4), said first display area to display one or more first symbols that represent said first server (page 1; page 3; page 4), said second server (page 1; page 3; page 4), or said first storage unit (page 1; page 3; page 4), and having second symbols that represent communication paths (page 1; page 3; page 4), said second display area having third symbols

that represent said second storage unit (page 1; page 3; page 4), wherein said management computer (page 1; page 3; page 4) selects the second storage unit (page 1; page 3; page 4) in accordance with receiving an indication for moving (page 1; page 3; page 4) one of said third symbols from said second display area into said first display area (page 1; page 3; page 4).

It would have been obvious to a person of ordinary skill in the art at the time of Applicant's invention to modify the teachings of Peters as modified by Baru with the teachings of Microsoft to include that the management computer includes a display unit operable to present a first display area and a second display area, said first display area to display one or more first symbols that represent said first server, said second server, or said first storage unit, and having second symbols that represent communication paths, said second display area having third symbols that represent said second storage unit, wherein said management computer selects the second storage unit in accordance with receiving an indication for moving one of said third symbols from said second display area into said first display area with the motivation to move files from and folders from one folder to another (Microsoft, page 4).

Conclusion

8. **THIS ACTION IS MADE FINAL.** See MPEP § 706.07(a). Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within

TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the date of this final action.

Points of Contact

9. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Apu M. Mofiz whose telephone number is (571) 272-4080. The examiner can normally be reached on Monday – Thursday 8:00 A.M. to 4:30 P.M.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Jeffrey Gaffin can be reached at (571) 272-4146. The fax numbers for the group is (571) 273-8300.

Any inquiry of a general nature or relating to the status of this application should be directed to the Group receptionist whose telephone number is (703) 305-9600.

Apu M. Mofiz
Primary Patent Examiner
Technology Center 2100

July 26, 2005